

A QMA-Complete Translationally Invariant Hamiltonian Problem and the Complexity of Finding Ground State Energies in Physical Systems

Alastair Kay

*Centre for Quantum Computation, DAMTP, Centre for Mathematical Sciences,
University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK*

(Dated: February 1, 2008)

Here we present a problem related to the local Hamiltonian problem (identifying whether the ground state energy falls within one of two ranges) which is restricted to being translationally invariant. We prove that for problems with a fixed local dimension and $O(\log(N))$ -body local terms, or local dimension N and 2-body terms, there are instances of the problem which are QMA-complete. We discuss the implications for the computational complexity of finding ground states of these systems, and hence for any classical approximation techniques that one could apply including DMRG, Matrix Product States and MERA. One important example is a 1D lattice of bosons with nearest-neighbor hopping at constant filling fraction i.e. a generalization of the Bose-Hubbard model.

Introduction: The complexity class QMA is the quantum analogue of NP i.e. it is the class of problems for which there exists a polynomially sized quantum computation V that can verify the solutions to the problem, outputting ‘yes’ if the input is indeed a solution, and ‘no’ otherwise. Recent interest has focused around establishing complete problems for QMA, as these encapsulate the character of the class. This started with Kitaev’s proof that $O(\log N)$ -local and 5-local Hamiltonian problems are QMA-complete [1], where a Hamiltonian problem is defined to be establishing whether a particular Hamiltonian has its smallest eigenvalue less than one threshold, a , or greater than another, b , where $b > a$. Since this original proof, the structure of Hamiltonians which still solve QMA-complete problems has been reduced to 2-local Hamiltonians [2], and their spatial organisation in 2D [3] and 1D lattices [4] has also been demonstrated. In the one-dimensional case, this was achieved by increasing the local dimension of spins from 2 to 12. Reductions of other, seemingly unrelated, problems have also been shown, hence proving their QMA-completeness, such as the N -representability problem [5].

In this paper, we impose the translational invariance over a one-dimensional (periodic) chain on the Hamiltonian, which can elucidate the trade-offs in resources required for quantum computation. One extreme that we already know is given by the aforementioned local Hamiltonian constructions [6], where we have no time resolution in our control of the system, only spatial resolution. Another extreme is where we exchange the spatial control for temporal control [7], which is the culmination of the study of global control schemes (see [8] and references therein). The problem that we develop here, attempting to find a translationally invariant Hamiltonian for computation, has neither spatial nor temporal control. Consequently, there is a cost elsewhere in the scheme - either in the local dimension of the spins used ($\sim \text{poly}(N)$ where there are N qubits), or the range of the Hamiltonian terms ($\sim \log(N)$).

The ultimate aim of this study, however, is to examine the computational complexity of finding the ground state energy of a translationally invariant local Hamil-

tonian. It is typically argued that since there are exponentially many parameters required to define a state, determining the ground state of an arbitrary Hamiltonian is a hard problem. However, by restricting the symmetry of a system, the effective number of parameters is reduced and thus one might hope to efficiently approximate the ground state of, say, translationally invariant local Hamiltonians. This is a hugely important task in condensed matter physics since these are the Hamiltonians that typically occur in nature, and is precisely the problem that is tackled by a range of approximation techniques such as DMRG/Matrix Product States (MPS) [9, 10, 11], Quantum Monte Carlo simulations [12] and MERA [13]. There are good reasons to expect that efficient representations of these states exist [14, 15], and so it becomes a question of the computational difficulty of finding a good representation. One strikingly open question is “how accurately can the ground state of a one-dimensional translationally invariant Hamiltonian and the corresponding energy be approximated given polynomial time?”, and it is hoped that one could discover techniques where it is possible to present certificates that guarantee a minimum level of accuracy. Some steps have been taken in understanding this problem, providing NP-hard instances of MPS [16], and several classifications for the two-dimensional generalisation [17, 18]. Developing QMA-complete instances of the translationally invariant local Hamiltonian problem will show that no approximation strategy (that is applicable to these cases) can give a certificate without first resolving the question $\text{QMA} = \text{P}$? This is in contrast to the previous discussions [16, 17, 18] which are tied to specific approximation strategies.

Definitions: For completeness, we repeat the following definitions:

Definition 1. (QMA) A promise problem L with ‘yes’ instances L_{yes} and ‘no’ instances L_{no} is in QMA if there exists a quantum polynomial time verifier V and polynomial p such that for some $\varepsilon(M) = 2^{-\tilde{\Omega}(M)}$

- $\forall x \in L_{\text{yes}}, \exists |\xi\rangle \quad \Pr(V(|x\rangle, |\xi\rangle) = 1) \geq 1 - \varepsilon$
- $\forall x \in L_{\text{no}}, \forall |\xi\rangle \quad \Pr(V(|x\rangle, |\xi\rangle) = 1) \leq \varepsilon$

where $|x| = M$, $|\xi\rangle \in \mathcal{B}^{\otimes p(M)}$ and $\Pr(V(|x\rangle, |\xi\rangle) = 1)$ denotes the probability that V outputs 1 given $|x\rangle$ and $|\xi\rangle$.

Definition 2. Given a 2-local translationally invariant Hamiltonian on a one-dimensional chain of N D -dimensional systems, $H = \sum_{j=1}^N h_{j,j+1}$ with $\|h\| \leq \text{poly}(N)$, and two constants $a < b$, the promise problem that we examine has yes instances in which the smallest eigenvalue of H is at most a , and no instances in which it is larger than b , and we must decide which is the case. We call this problem TI2LH (Translationally Invariant 2-Local Hamiltonian).

Computation from a Hamiltonian: The central part of previous proofs of QMA-completeness, as in this one, is the construction of Hamiltonians by showing how to implement an arbitrary quantum computation within the Hamiltonian (and hence can implement the verifier V). We shall assume that this computation acts on N qubits, of which the first M contain the string x to be verified, and the other $N - M$ are initialised in $|0\rangle$, and are used as ancillas in the computation. This computation is typically achieved by introducing a ‘clock’ system, and a Hamiltonian that implements a particular gate as the clock increments. The Hamiltonian hops the system through the different clock states, in a way that maps directly to quantum state transfer along a spin chain [19, 20, 21]. In the original proposal [1], the hopping terms between neighboring clock states all had the same strengths. Since a uniformly coupled chain does not achieve perfect transfer of a state through the system [19, 21] (and hence perfect arrival of the computation in the final state is impossible), the computation can be extended by a number of identity operations, thus ensuring that the probability that the computation has got past the final step is high [22]. It has since been observed [6] by making an analogy with perfect state transfer schemes [20, 21] that varying the coupling strengths can allow perfect arrival of the computation, and hence the addition of further steps is unnecessary. However, here we will continue to use the uniformly coupled case due to the simplicity of the related eigenvectors [23], particularly since for the proof of QMA-completeness, we do not actually perform the computation, and hence do not need perfect arrival of the result.

Ultimately, to implement a computation composed of a discrete set of gates, we need to introduce a clock system. However, in a translationally invariant Hamiltonian, we neither have time resolution nor spatial resolution in which to encode this clock. Instead, we choose to hold it in the state of each spin. Each of these D -dimensional spins will have a number of states: $|on\rangle$, and $|x, y, z\rangle$, where $x \in \{0, 1\}$ is the qubit state which contains the computational component of the state, $z \in \{1, 2, \dots, N\}$ denotes the position of the spin in the lattice and $y \in \{0, 1, \dots, R\}$ is the time component, which we think of as ‘cycle number’. The input state to our computation will be a state of $N + 1$ spins, $|on\rangle \bigotimes_{n=1}^N |x_n, n, 0\rangle$

where x_n are the bits of the string x to be tested, and the $|on\rangle$ state indicates the initial position for a ‘read/write-head’ that our translationally invariant Hamiltonian will be able to scan backwards and forwards across the chain in discrete cycles (it will always scan across all N computational qubits). As it moves across, the cycle number is locally incremented by 1, and the required local unitaries can be implemented due to the position label. The computing part of the Hamiltonian takes the form

$$\begin{aligned} h_{i,j}^{\text{comp}} = & -|on\rangle\langle on|_i \otimes |0, 1, 0\rangle\langle 0, 1, 1|_j \\ & -U_{n,m}^{i,j} \otimes |m, n\rangle\langle m - \delta_m, n|_i \otimes |m, n'\rangle\langle m - 1 + \delta_m, n'|_j \\ & -U_{N,m}^{i,j} \otimes |m, N - 1\rangle\langle m, N - 1|_i \otimes |m + 1, N\rangle\langle m - 1, N|_j \\ & -U_{1,m}^{i,j} \otimes |m + 1, 1\rangle\langle m - 1, 1|_i \otimes |m, 2\rangle\langle m, 2|_j + \text{h.c.} \end{aligned}$$

where we are implicitly summing over cycle index m and position index n ($n' = n + 1$, $2\delta_m = 1 + (-1)^m$). The unitaries $U_{n,m}^{i,j}$ implement the step of the computation that is required to be acted at time m between qubits labelled n and $n + 1$, and only acts on the qubit space of sites i and j . The first term initialises the read/write head. The next term propagates the head to the right or left depending on the value of δ , and hence the parity of the cycle number. The final terms are responsible for reversing the direction of travel of the head when it reaches either extreme of the chain. For convenience, we fix the ground state energy of this Hamiltonian to zero by also adding the terms

$$2\mathbf{1}^i \otimes |m, n\rangle\langle m, n|_i - \mathbf{1}^i \otimes (|0, 1\rangle\langle 0, 1| + |R, 1\rangle\langle R, 1|)_i.$$

The final two terms mark a departure from the correspondence to transfer on a uniformly coupled chain, but makes the form of the lowest energy eigenvector of the Hamiltonian much simpler (see Eqn. (1)).

We now have a translationally invariant Hamiltonian $H_{\text{comp}} = \sum_i h_{i,i+1}$ where the terms $\|h\| \leq \text{poly}(N)$. However, it currently acts on a non-translationally invariant input state. It is helpful to observe that the input states $|\eta^n\rangle$, where $|on\rangle$ is located on qubit n can be combined to make the state translationally invariant, $\sum_n |\eta^n\rangle$. After computation, the result would still be available because projective measurement would locate the position of $|on\rangle$. Alternatively, one could use the elegant techniques described in [7] to achieve computation on a translationally invariant input state.

QMA-Completeness: Now that we have a translationally invariant Hamiltonian for performing computations, we turn our attention to proving the QMA-completeness of TI2LH. We do this by adding some extra terms to the Hamiltonian

$$H = J_1 H_{\text{input}} + J_2 (\alpha H_{\text{form}} + H_{\text{comp}}) + R(N - 1) H_{\text{output}}.$$

Intuitively, the terms H_{input} and H_{form} will verify that we have the desired input by adding an energy penalty for incorrect states. H_{input} is particularly used to ensure the state of ancillas for the computation. We take these

to be positioned on qubits $M + 1$ to N , and must be in the $|0\rangle$ state at time 0. Thus, we require an energy penalty

$$h_i^{input} = \sum_{n=M+1}^N |1, 0, n\rangle \langle 1, 0, n|_i.$$

The terms in H_{form} , while they could be applied just when the clock is at times 0, can equally well be applied at all time because they act on parts of the state that don't change, such as the $|on\rangle$ and position states. As such, they are more conveniently considered in conjunction with H_{comp} . The first term we apply is to ensure that there is at least one $|on\rangle$ present in the system, $-|on\rangle \langle on|_i$. However, we don't want more than one of these states, so we also need an energy cost for an $|on\rangle$ being on the right of a state in position $n < N$, that more than compensates for the reduced energy due to the $|on\rangle \langle on|$ term,

$$\sum_{n < N} 2|x, m, n\rangle \langle x, m, n|_i \otimes |on\rangle \langle on|_{i+1}.$$

Similarly, we should ensure that neighboring position labels only increment by 1,

$$\sum_{q \neq 1} |x, m, n\rangle \langle x, m, n|_i \otimes |x, m, n+q\rangle \langle x, m, n+q|_{i+1}.$$

Equally, H_{output} will test the final states, and add an energy penalty for a ‘no’ result,

$$h_i^{output} = |1, R, 1\rangle \langle 1, R, 1|_i,$$

where the computation outputs the accept/reject code ($|0\rangle / |1\rangle$) on the first qubit after $R = \text{poly}(N)$ cycles.

We must now verify that the eigenstates of H satisfy the required properties of TI2LH. We choose to closely follow the proof presented in [2], making use of their projection lemma,

Lemma 1. *Let $H = H_1 + H_2$ be the sum of two Hamiltonians operating on some Hilbert space $\mathcal{H} = \mathcal{S} + \mathcal{S}^\perp$. The Hamiltonian H_2 is such that \mathcal{S} is a zero eigenspace and the eigenvectors in \mathcal{S}^\perp have eigenvalue at least $J > 2\|H_1\|$. Then,*

$$\lambda(H_1|_{\mathcal{S}}) - \frac{\|H_1\|^2}{J - 2\|H_1\|} \leq \lambda(H) \leq \lambda(H_1|_{\mathcal{S}}).$$

$\lambda(H)$ denotes the smallest eigenvalue of H .

In particular, one can select, say $J = 8\|H_1\|^2 + 2\|H_1\|$ in order to provide the bounds

$$\lambda(H_1|_{\mathcal{S}}) - \frac{1}{8} \leq \lambda(H) \leq \lambda(H_1|_{\mathcal{S}}).$$

Firstly, let's consider that there are ‘yes’ instances x with bits x_n (and $x_n = 0$ if $n > M$). To test this solution, the input to the computation would have been

$$|\eta_0\rangle = |on\rangle \bigotimes_{n=1}^N |x_n, 0, n\rangle.$$

We can consider the $(N-1)R$ discrete steps of the computation and write the state after each step as $|\eta_n\rangle$. Using these, we can write down a state

$$|\eta\rangle = \frac{1}{\sqrt{R(N-1)+1}} \sum_n |\eta_n\rangle \quad (1)$$

and evaluate that

$$\langle \eta | H_{input} | \eta \rangle = -1, \quad \langle \eta | H_{comp} | \eta \rangle = 0.$$

Therefore, if the verifier accepts with probability $\geq 1 - \varepsilon$, $\langle \eta | H | \eta \rangle \leq -J_1 + \varepsilon$.

Now we must prove that if there are no ‘yes’ instances, all eigenvalues are larger than some value b . One can precisely follow the arguments of [2]. However, there is one small additional detail that we must first consider – the energy gap between the ground-state space of $\alpha H_{form} + H_{comp}$ and the next excited state. Firstly we observe that due to H_{form} , the space splits into a series subspaces defined by the total number of $|on\rangle$ states present, and how badly ordered the position labels are. All of these subspaces have an offset of at least α from the original subspace of a single $|on\rangle$ state and correctly ordered position labels. Within this original subspace, we have the ground state subspace described by states $|\eta\rangle$ for all possible input values x , and cyclic permutations of the initial position of $|on\rangle$. Relative to the ground state space, all eigenvectors are separated by at least $c/(R(N-1))^2$ for some positive constant c [1]. This is readily proved by performing a transformation to the state transfer model. Thus, by setting $\alpha > c/(R(N-1))^2$, we can ensure that the minimum energy gap is at least $c/(R(N-1))^2$. We can now proceed with application of the projection lemma where $H_2 = J_2(\alpha H_{form} + H_{comp})$ and $H_1 = H - H_2$. As in [2], we need to nest the application of this lemma, so that $H' = H_1|_{\mathcal{S}^{comp}}$, $H'_2 = J_1 H_{input}|_{\mathcal{S}^{comp}}$ and $H'_1 = H' - H'_2$ before finally arriving at the result that $\lambda(H) \geq 3/4 - \varepsilon$, thereby proving that solving the local Hamiltonian problem implies the ability to solve all QMA problems.

We can transform this Hamiltonian of $N + 1$ spins of dimension $D = 2RN + 1$ made up of two-body terms into a Hamiltonian acting on $N \log_2(D-1) + 1$ copies of 7-level spins and $2 \log_2(D-1)$ -body terms. This transformation can be achieved by replacing each D -dimensional spin with a set of spins where one level is used as the $|on\rangle$ state, two levels encode a qubit state and the final three levels ($3 \times 2 + 1 = 7$) denote whether that spin is being used to store a computational qubit, or a position or a cycle label, where an entire position and cycle label is achieved by combining $\log_2(D-1)$ neighbors. We require N of these blocks, and a single extra spin to be initialised in the $|on\rangle$ state. The two-body terms from H^{comp} need to act across two complete blocks of $\log_2(D-1)$ spins.

Consequences and Conclusions: In summary, if there exist ‘yes’ instances to a QMA problem, the minimum eigenvalue of the corresponding H (implementing the verifier V) is $\leq -J_1 + \varepsilon$, whereas if there are no such instances, the minimum eigenvalue is at least $-J_1 + \frac{3}{4} - \varepsilon$.

Therefore, solving TI2LH allows you to solve any QMA problem. Note that our construction of the Hamiltonian is necessarily degenerate [24], as seen by considering the case where there is a ‘yes’ solution, tested by $|\eta\rangle$. There are $N+1$ equivalent states corresponding to the different possible starting positions of $|on\rangle$, all of the same energy. A similar argument will hold for the case where there are no such solutions. Our construction also implicitly shows how one can take a system on N spins of dimension d with k -body local Hamiltonian terms and make it into a Hamiltonian acting on N spins of dimension Nd with k -body local terms which is translationally invariant while preserving the structure of the eigenstates.

This proves that any attempt to classically approximate the ground state energy of a degenerate Hamiltonian of N spins where either the Hamiltonian is composed of $O(\log(N))$ -body (local) terms on spins of fixed dimension, or 2-body terms acting on spins of $\text{poly}(N)$ dimension is QMA-hard (in N) and thus, assuming $\text{QMA} \neq \text{P}$, will, in the worst case, be intractable. Specifically, by polynomially increasing the computational resources available for determining the ground state energy, the accuracy, in the worst case, cannot be more than polynomially increased. Falling within this class of Hamiltonians is the Bose-Hubbard model at constant filling fraction. At a filling fraction of 1, for example, if there are N lattice sites, then there are N bosons. We have to allow for the possibility that all N bosons could be found in a single lattice site, and thus each lattice site must be treated as an N -level system.

As we have proven, while there is a degenerate ground state, there are also excited states of these Hamiltonians. Since there is an energy gap, properties such as exponentially decaying correlation functions [15] hold, and techniques such as DMRG/MPS, and Vidal’s re-

cently proposed Multi-scale Entanglement Renormalization Ansatz (MERA) [13] provide an efficient representation. We conclude that finding this representation must be a computationally hard task. This does not say, of course, that classical approximation strategies can’t be useful in a range of physically important problems since the QMA-hardness of a problem only elucidates the worst case scenario. In contrast to previous attempts at similar problems [16, 17, 18], the example Hamiltonian that we have produced is not tied to a specific algorithm that one might be trying to implement, but instead applies to the general question that we are trying to solve with such an algorithm, thereby eliminating the possibility of globally efficient algorithms that encapsulate the Hamiltonians in question.

Evidently, an important question for the future is whether these results can be extended to the situation where both the local dimension and the Hamiltonian terms remain fixed as the system size grows, since this is the case which is most closely related to physical systems. While it seems unlikely that our technique can be directly applied (as we need to use a clock), there are strong indications that this may well be the case. In the first case we analysed, where the local Hilbert space dimension scales with $\text{poly}(N)$, but with two-body terms in the Hamiltonian, we know most of the structure of the ground state from Eqn. (1) if there are ‘yes’ solutions. By imposing this structure on the ansatz states of our approximation strategy, we are effectively reduced to minimising the energy over the local qubit states i.e. the QMA-hardness of this problem is effectively encoded in a nearest-neighbor Hamiltonian acting on qubits.

We would like to thank Lluis Masanes and Jens Eisert for useful conversations. This work was supported by Clare College, Cambridge.

-
- [1] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, Graduate studies in mathematics (American Mathematical Society, Providence, Rhodes Island, 2002).
 - [2] J. Kempe, A. Kitaev, and O. Regev, SIAM Journal of Computing **35**, 1070 (2006).
 - [3] R. Oliveira and B. M. Terhal (2005), quant-ph/0504050.
 - [4] D. Aharonov, D. Gottesman and J. Kempe (2007), arXiv: 0705.4077; S. Irani (2007), arXiv: 0705.4067.
 - [5] Y. K. Liu, M. Christandl, and F. Verstraete, Phys. Rev. Lett. **98**, 110503 (2007).
 - [6] A. Kay, quant-ph/0702092.
 - [7] K. G. H. Vollbrecht and J. I. Cirac, Phys. Rev. A **73**, 012324 (2006).
 - [8] A. Kay (2007), quant-ph/0702239.
 - [9] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
 - [10] S. R. White, Phys. Rev. B **48**, 10345 (1993).
 - [11] F. Verstraete, D. Porras, and J. I. Cirac, Phys. Rev. Lett. **93**, 227205 (2004).
 - [12] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Rev. Mod. Phys. **73**, 33 (2001).
 - [13] G. Vidal (2006), quant-ph/0610099.
 - [14] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).
 - [15] M. B. Hastings, Phys. Rev. B **73**, 085115 (2006).
 - [16] J. Eisert, Phys. Rev. Lett. **97**, 260501 (2006).
 - [17] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **98**, 140506 (2007).
 - [18] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Physical Review Letters **96**, 220601 (2006).
 - [19] S. Bose, Phys. Rev. Lett. **91**, 207901 (2003).
 - [20] M. Christandl, N. Datta, A. Ekert, and A. J. Landahl, Phys. Rev. Lett. **92**, 187902 (2004).
 - [21] M. Christandl *et al.*, Phys. Rev. A **71**, 032312 (2005).
 - [22] D. Aharonov *et al.*, FOCS **00**, 42 (2004).
 - [23] In particular, the lowest energy eigenvector of an N -qubit chain takes the form $\frac{1}{2^{N/2}} \sum_n (-1)^n \sqrt{\binom{N}{n}} |n\rangle$, and hence the final coefficient is exponentially suppressed
 - [24] If we use open boundary conditions instead of periodic, the degeneracy due to permutation invariance is lifted.